

(21) Application No 9204329.8

(22) Date of filing 28.02.1992

(71) Applicant
Sony United Kingdom Limited

(Incorporated in the United Kingdom)

Sony House, South Street, Staines,
Middlesex, TW18 4PF, United Kingdom

(72) Inventor
Clive Henry Gillard

(74) Agent and/or Address for Service
D Young & Co
21 New Fetter Lane, London, EC4A 1DA,
United Kingdom

(51) INT CL⁵
H04N 11/02

(52) UK CL (Edition L)
H4F FD1B1 FD1D1 FD3 FRX

(56) Documents cited
None

(58) Field of search
UK CL (Edition K) H4F FRC FRD FRG FRM FRP
FRR FRT
INT CL⁵ H04N
Online databases: WPI

(54) Image data compression

(57) Data compression apparatus for compressing luminance Y and chrominance U/V portions of input image data to form compressed image data having an overall compression factor is described. The data compression apparatus comprises luminance compression means 12', 64, 18', 20' for compressing the luminance portion of the input image data by a luminance compression factor and chrominance compression means 12'', 66, 18'', 20'' for compressing the chrominance portion of the input image data by a chrominance compression factor. In response to a feedback signal 74 dependent on the overall compression factor a control means 76 coupled to the luminance compression means and the chrominance compression means controls the overall compression factor by varying the luminance compression factor, the chrominance compression factor and the ratio of the luminance compression factor to the chrominance compression factor.

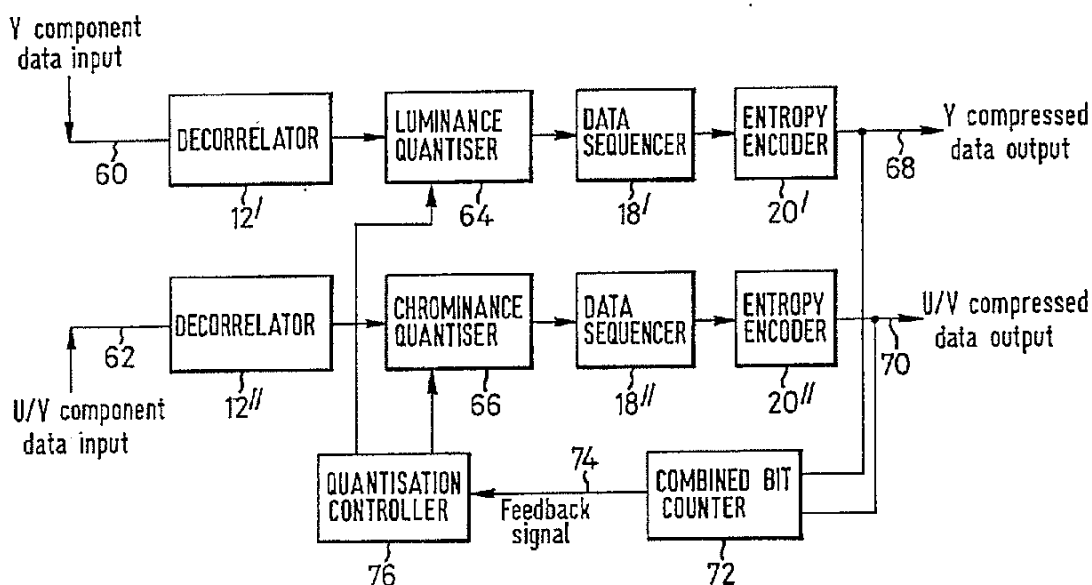


Fig.8

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

The print reflects an assignment of the application under the provisions of Section 30 of the Patents Act 1977.

GB 2 266 635 A





Fig.3

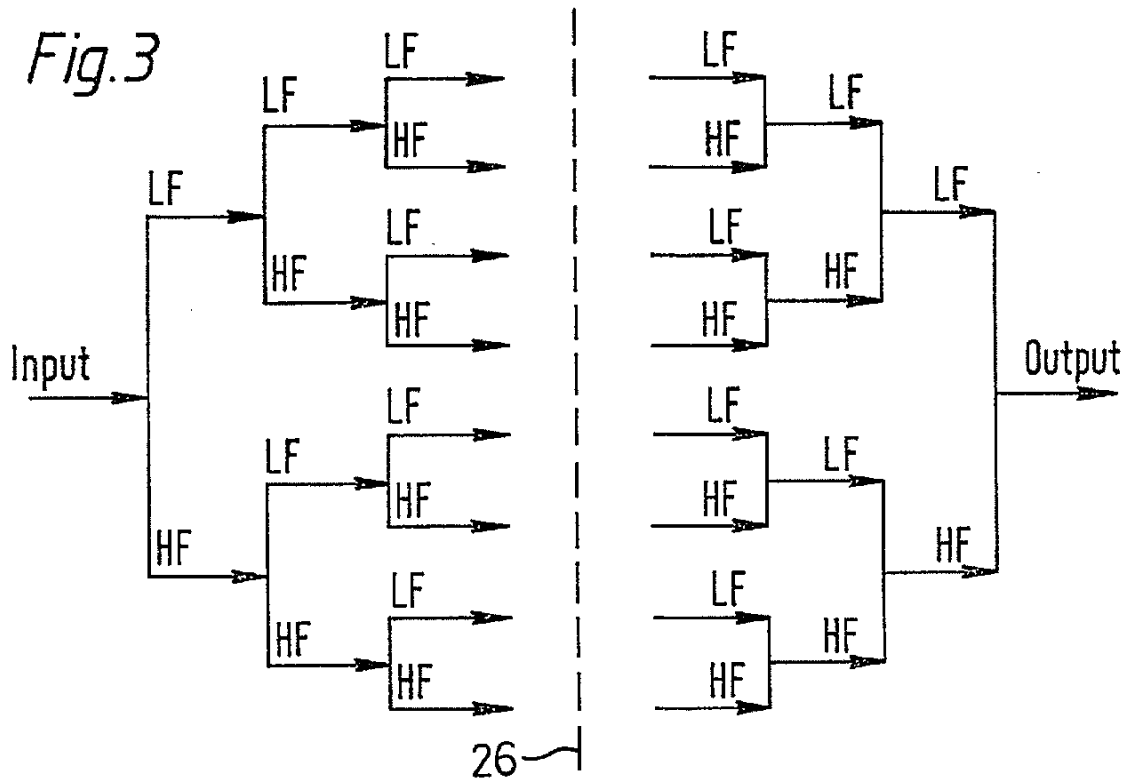


Fig.4






Fig. 5

4/8

α_{00}	α_{01}	α_{02}	α_{03}	α_{04}	α_{05}	α_{06}	α_{07}
α_{10}	α_{11}	α_{12}	α_{13}	α_{14}	α_{15}	α_{16}	α_{17}
α_{20}	α_{21}	α_{22}	α_{23}	α_{24}	α_{25}	α_{26}	α_{27}
α_{30}	α_{31}	α_{32}	α_{33}	α_{34}	α_{35}	α_{36}	α_{37}
α_{40}	α_{41}	α_{42}	α_{43}	α_{44}	α_{45}	α_{46}	α_{47}
α_{50}	α_{51}	α_{52}	α_{53}	α_{54}	α_{55}	α_{56}	α_{57}
α_{60}	α_{61}	α_{62}	α_{63}	α_{64}	α_{65}	α_{66}	α_{67}
α_{70}	α_{71}	α_{72}	α_{73}	α_{74}	α_{75}	α_{76}	α_{77}

Fig.7




Fig.9




Fig.8

Fig. 10





FIG. 11



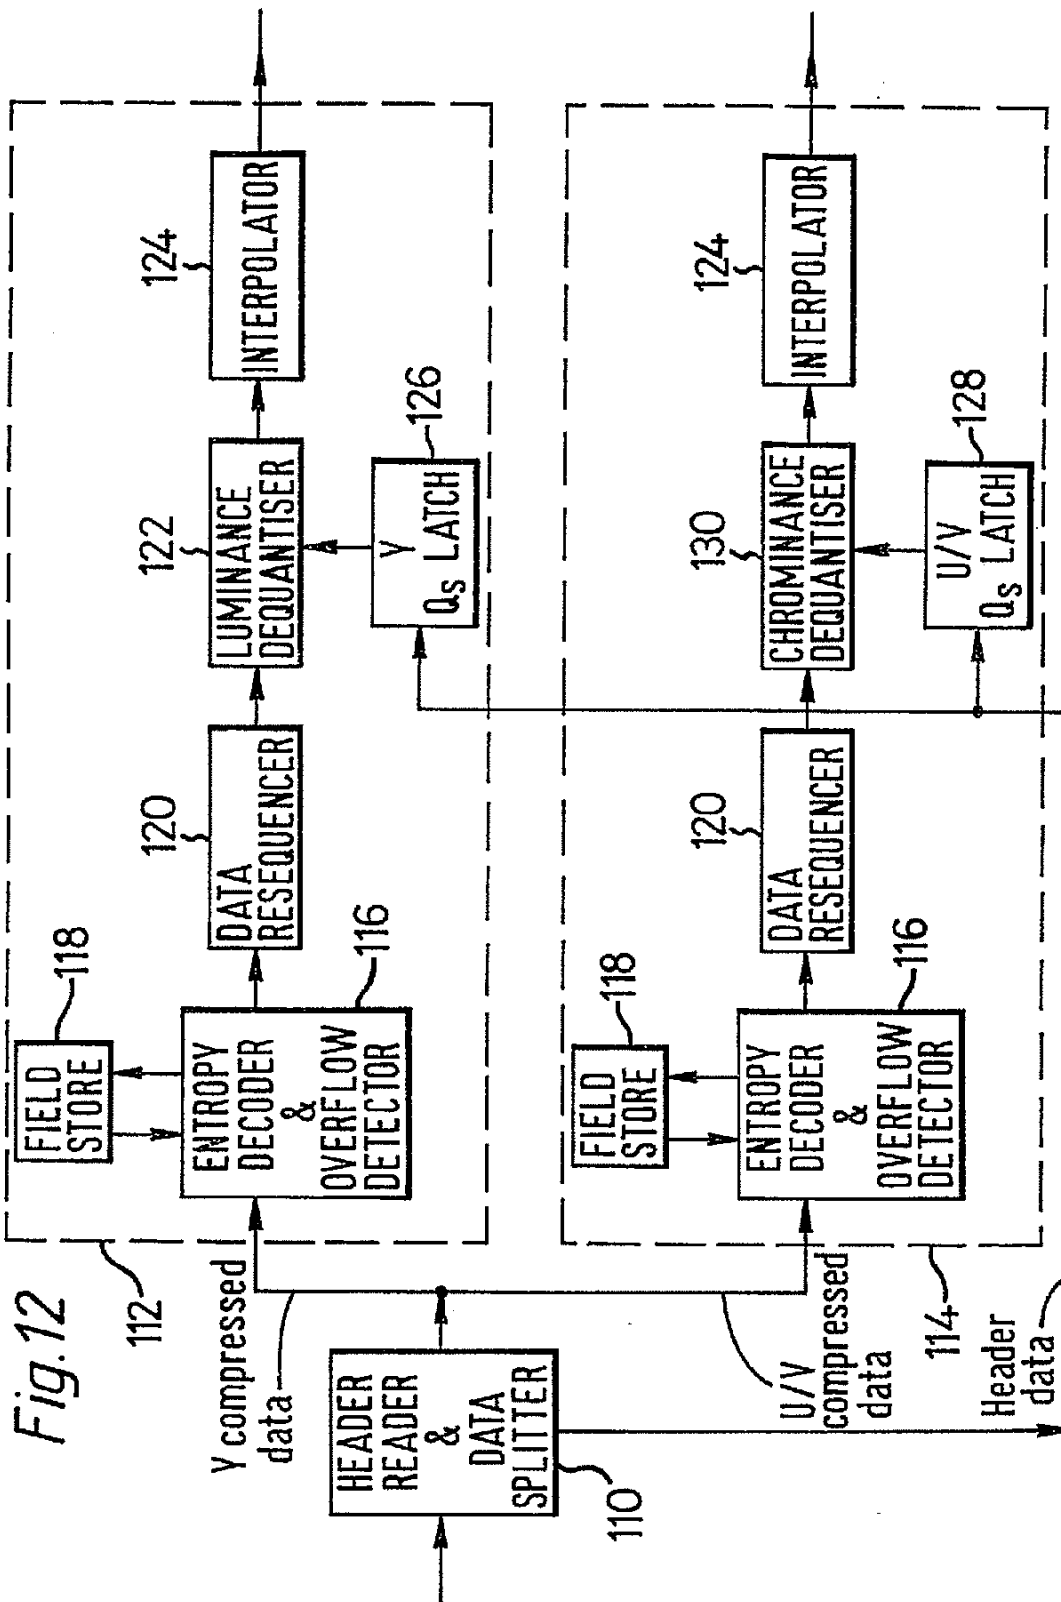


IMAGE DATA COMPRESSION

This invention relates to the field of image data compression.

5 In some known image data compression systems a quantiser is used
to selectively remove information from luminance and chrominance
components of input image data as part of the compression process. A
higher degree of quantisation will produce a higher degree of
compression. In general, it is not possible to determine in advance
10 what degree of compression will be produced by a particular degree of
quantisation, since different portions of the input image data may vary
in information content and suitability for compression.

In some systems this is not a problem. For example, in the image
data compression system proposed by the Joint Photographic Experts
Group and currently under review by the International Standards
15 Organisation input image data is compressed using a quantiser and
entropy encoder to a degree that varies upon the particular image and
parameters used. The compressed image data is intended for non-real
time processing and display and, accordingly, variations in the degree
of compression achieved do not cause particular problems.

20 In contrast to the above, the problem of variation in the degree
of compression achieved can be particularly serious when compressing
input image data to occupy data blocks of a fixed size or to comply
with a similar constraint on the size of the compressed data. The
fixed size of the blocks of compressed image data may be the result of
25 the need to store each block on magnetic tracks of a fixed maximum
capacity or the need to use the blocks of compressed data in a real
time transmission system with a fixed maximum time for the processing
of each block. In such systems there is no flexibility in the maximum
size of blocks of compressed image data that can be accommodated. If
30 too much compressed image data is produced to fit into the fixed size
blocks, then an overrun with a consequential uncontrolled loss of data
will occur.

On the other hand, it is desirable not to compress the input
image data more than is necessary, in order to reduce any possible
35 deterioration of an image reconstructed from the compressed image data
which might be caused by the data compression processing.

A data compression system proposed in the copending British

Patent Application number 9119985.1 provides a dynamic quantisation control in order to control the degree of compression applied by the data compression system on a block-by-block basis.

5 It is a constant aim of image data compression systems to reduce the subjective impact of image data compression on the reconstructed image, while still providing the required degree of compression.

10 This invention provides data compression apparatus for compressing luminance and chrominance portions of input image data to form compressed image data having an overall compression factor, the data compression apparatus comprising: luminance compression means for compressing the luminance portion of the input image data by a luminance compression factor; chrominance compression means for compressing the chrominance portion of the input image data by a chrominance compression factor; feedback means for supplying a feedback
15 signal dependent on the overall compression factor; and control means responsive to the feedback signal and coupled to the luminance compression means and the chrominance compression means for controlling the overall compression factor by varying the luminance compression factor, the chrominance compression factor and the ratio of the
20 luminance compression factor to the chrominance compression factor.

The invention recognises that the subjective impact of variations in luminance image data compression can be different to that of equal variations in chrominance image data compression. This fact is exploited in the invention in which the luminance compression factor,
25 the chrominance compression factor and the ratio of the luminance and chrominance compression factors can be altered in order to vary the overall compression factor of the image data compression apparatus. Thus the luminance compression factor and the chrominance compression factor used to achieve a particular value of the overall compression factor can be related to one another according to a predetermined law,
30 the aim of which can be to reduce the subjective impact of the compression on the reconstructed image.

The term 'compression factor' has been used in the sense that if the compression factor is high, then the overall degree of compression
35 is high and the resulting quantity of compressed data will be small. If the compression factor is low, then the overall degree of compression is low and the resulting quantity of compressed data will

be large.

It will be appreciated that by defining two of the three parameters mentioned above (the luminance compression factor, the chrominance compression factor and the ratio of the luminance compression factor to the chrominance compression factor), the remaining parameter is defined.

In a preferred embodiment the control means varies the luminance compression factor to be less than or equal to the chrominance compression factor. In particular, it is preferred that the control means varies the luminance compression factor to be equal to the chrominance compression factor for one or more predetermined ranges of values of the overall compression factor (such as very large and very small values of the overall compression factor) and to be less than the chrominance compression factor for values of the overall compression factor which are not in the one or more predetermined ranges of values. In this arrangement, given a limit on the quantity of compressed data which can be produced, the quality of the chrominance data can be sacrificed by increasing the chrominance compression factor, in order to release more capacity for the subjectively more important luminance data.

In a preferred embodiment the luminance compression means and the chrominance compression means each comprise: a decorrelator for transforming the image data into transformed data representing a plurality of frequency components in the spatial frequency domain; a quantiser for applying a variable degree of quantisation to the transformed data; and an entropy encoder for encoding the quantised transformed data.

In accordance with the invention a variable degree of compression is applied to the luminance and chrominance portions of the input image data and consequently a variable degree of decompression must subsequently be applied to the compressed luminance and chrominance data. In order to control this action preferred embodiments of the invention provide means for associating a data header with the compressed image data indicative of the luminance compression factor and the chrominance compression factor corresponding to that compressed image data.

Preferably, in a case in which blocks of the input image data are

compressed to form blocks of compressed data, the feedback means comprises a compressed bit counter for storing a count of bits in blocks of compressed data. For blocks of input image data of a known size, this bit count will be directly related to the overall compression factor.

5 Viewed from a second aspect this invention provides data decompression apparatus for decompressing compressed image data having luminance and chrominance portions, the data decompression apparatus comprising: header reading means for reading a data header associated with the compressed image data and providing separate indications of the degree of luminance compression and the degree of chrominance compression applied to the luminance and the chrominance portions respectively of the compressed image data; luminance decompression means, responsive to the header reading means, for decompressing the luminance portion of the compressed image data by a degree dependent on the degree of luminance compression; and chrominance decompression means, responsive to the header reading means, for decompressing the chrominance portion of the compressed image data by a degree dependent on the degree of chrominance compression.

10 Viewed from a third aspect this invention provides a method of compressing luminance and chrominance portions of input image data to form compressed image data having an overall compression factor, the method comprising the steps of: compressing the luminance portion of the input image data by a luminance compression factor; compressing the chrominance portion of the input image data by a chrominance compression factor; and varying the luminance compression factor, the chrominance compression factor and the ratio of the luminance compression factor to the chrominance compression factor in order to control the overall compression factor.

15 Viewed from a fourth aspect this invention provides a method of decompressing compressed image data having luminance and chrominance portions, the method comprising the steps of: reading a data header associated with the compressed image data and providing separate indications of the degree of luminance compression and the degree of chrominance compression applied to the luminance and the chrominance portions respectively of the compressed image data; decompressing the luminance portion of the compressed image data by a degree dependent on

the degree of luminance compression; and decompressing the chrominance portion of the compressed image data by a degree dependent on the degree of chrominance compression.

5 The invention may be particularly usefully employed in image data storage apparatus.

Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings, throughout which like parts are referred to by like references, and in which:

10 Figure 1 is a schematic illustration of a data compression system;

Figure 2 is a schematic illustration of one stage of sub band coding;

Figure 3 is a schematic illustration of a higher order sub band coding system;

15 Figure 4 illustrates a two-dimensional sub band decorrelator;

Figure 5 illustrates a frequency separated video signal;

Figure 6 illustrates the human psychovisual response to images of differing spatial frequency;

Figure 7 is a quantisation matrix;

20 Figure 8 is a schematic diagram of a data compression apparatus providing separate luminance and chrominance compression;

Figure 9 is a graph showing the variation of the luminance and chrominance compression factors in the apparatus of Figure 8;

25 Figure 10 is a more detailed schematic diagram of a data compression apparatus providing separate luminance and chrominance compression;

Figure 11 shows a technique for determining changes to the luminance and chrominance compression factors in the apparatus of Figure 10; and

30 Figure 12 shows a decompression apparatus complementary to the compression apparatus of Figure 10.

Figure 1 shows an apparatus for effecting intra-image frequency separation and compression of a video signal in the two-dimensional spatial frequency domain. A video signal, which is in digital form and comprises successive multi-bit (for example 8-bit) samples or words each representing a respective pixel of the luminance or chrominance component of a scanned image or picture, is applied via an input 10 to

35

a decorrelator 12. A frequency separated video signal is fed by the decorrelator 12 to a quantiser 14 and then via a data sequencer 18 to an entropy encoder 20, which together compress the frequency separated video signal provided by the decorrelator 12 to produce a compressed signal on an output 19. The compressed signal can then be transmitted or stored. After transmission or storage, the compressed signal can be restored substantially to its original form by expansion by way of entropy decoding, resequencing, dequantising and correlation operations which employ parameters converse to those used for decorrelation, sequencing, quantisation and entropy encoding; respectively, upon compression.

The operation of decorrelation performed in the decorrelator 12 relies upon the fact that neighbouring pixels of an image are highly correlated, whereby processing an image (for example, a field or frame of a video signal) to form frequency separated signal portions representing different components of the image in the two-dimensional spatial frequency domain enables a reduction in the amount of information needed to represent the image. Specifically, the frequency separated signal portions represent different spatial frequency components of the image.

The sequencing operation will be described in more detail below.

The quantisation operation, which is a lossy operation, in that it involves deliberate discarding of some frequency data considered to be redundant or of little importance to adequate perception of the image by the human psychovisual system, in itself enables some signal compression to be achieved. The quantiser 14 enables compression to be achieved in two ways: it reduces the number of levels to which the data input to it can be assigned, and it increases the probability of runs of zero value samples on the data it outputs. The ability to achieve enhanced signal compression provided by the operation of the quantiser is carried into effect in the entropy encoder 20 in that the reduction in information content achieved in the quantiser 14 enables a consequential bit (data) rate reduction to be achieved in the entropy encoder.

Further (non-lossy) compression, and bit (data) rate reduction, is provided in the entropy encoder 20 in which, in known manner, using for example variable length coding, the data produced by the quantiser

14 is encoded in such a manner that more probable (more frequently occurring) items of data produce shorter output bit sequences than less probable (less frequently occurring) ones. In this regard, the decorrelation operation has the effect of changing the probability distribution of the occurrence of any particular signal level, which is substantially the same as between the different possible levels before decorrelation, into a form in which it is much more probable that certain levels will occur than others.

The compression/coding system or apparatus as shown in Figure 1 can be embodied in a variety of ways, using different forms of decorrelation. An increasingly popular form of implementation makes use of so-called transform coding, and in particular the form of transform known as the discrete cosine transform. The use of discrete cosine transformation for decorrelation is in fact prescribed in a version of the compression system of Figure 1 described in the previously mentioned proposed standard prepared by the Joint Photographic Experts Group and currently under review by the International Standards Organisation. According to the transform technique of decorrelation, the signal is subjected to a linear transform (decorrelation) operation prior to quantisation and encoding. A disadvantage of the transform technique is that, although the whole image (for example, a whole field) should be transformed, this is impractical in view of the amount of data involved. The image (field) thus has to be divided into blocks (for example, of 8 x 8 samples representing respective pixels), each of which is transformed. That is, transform coding is complex and can be used on a block-by-block basis only.

A recently proposed approach to compression/coding in the frequency domain is that of sub-band coding. In this approach, the decorrelator 12 in the system of Figure 1 would comprise a spatial (two-dimensional) sub-band filtering arrangement which divides the input video signal into a plurality of uncorrelated sub-bands each containing the spatial frequency content of the image in a respective one of a plurality of areas of a two-dimensional frequency plane of the image, the sub-bands then being selectively quantised by the quantiser 14 in accordance with their positions in the sensitivity spectrum of the human psychovisual system. That is, decorrelation is achieved in

this case by putting the energy of the overall image into different sub-bands of the two-dimensional spatial frequency domain. Sub-band filtering is believed to provide better decorrelation than the transform approach. Also, unlike the transform technique, there is no
 5 restriction to operation on a block-by-block basis: the sub-band filtering can be applied directly to the video signal.

Figure 2 illustrates a sub-band coding system in which the input video signal is passed through a low-pass decimation filter 22 and a high-pass decimation filter 24. The resulting two output signals
 10 represent different portions of the frequency spectrum of the input signal. The two signals are then quantised, sequenced and entropy encoded as discussed in relation to Figure 1. The sub-band components of the input signal can now be transmitted or stored for later reproduction. The storage or transmission of the sub-band components
 15 is illustrated by the dashed line 26 in Figure 2.

When the sub-band components are recovered from the recording medium they are passed through corresponding matching filters to regenerate the original frequency components. These matching filters are a low-pass interpolation filter 30 and a high-pass interpolation
 20 filter 28. The outputs of the interpolation filters 28, 30 are added by a summation circuit 32 to yield the original video input signal.

Figure 2 illustrates the decomposition of the input video signal into two sub-bands. In practice, the input video signal would be decomposed into many more sub-band components. Figure 3 illustrates
 25 the decomposition of an input signal into eight sub-band components and its subsequent recombination into an output video signal.

The filters of the sub-band coding system comprise finite impulse response filters with appropriate delays and weighting coefficients to perform both horizontal and vertical frequency decomposition.
 30 Different forms of filter for performing sub-band frequency separation are known, e.g. some possible filters are described in the article entitled 'Exact Reconstruction Techniques for Tree Structured Sub-Band Coders', in IEEE Transactions on Acoustics, Speech and Signal Processing, Volume ASSP-34 at pages 434 to 441, June 1986.

35 Figure 4 illustrates the decorrelator 12 of Figure 1 in more detail. The decorrelator comprises a horizontal filter arrangement 46, an intermediate field store 48, a transpose sequencer (address

generator) 50, a vertical filter arrangement 52, an output field store 54 and an output sequencer (address generator) 56. Sub-band filtering is effected on a separable basis. Thus, in Figure 4, filtering in the two orthogonal image directions, namely the horizontal direction (the direction of image scanning in the case of conventional video) and the vertical direction, is effected entirely independently and separately of one another by respective one-dimensional filtering operations performed in the horizontal and vertical filter arrangements 46 and 52, respectively.

The horizontal filter arrangement 46 and vertical filter arrangement 52 can be of substantially the same construction as one another. Thus, the construction of the horizontal filter arrangement 46 only will be described in detail. The filtering is to achieve 8 sub-bands in each of the horizontal and vertical directions, that is to say that a square array of 64 (8×8) sub-bands is to be produced. The 64 sub-bands are to be of equal extent to one another.

The horizontal filter arrangement 46 is preferably of a tree or hierarchical structure as shown in Figure 3, comprising three successive filter stages. The first stage comprises a low pass filter (LF) and a high pass filter (HF), each of which is followed by a respective decimator (not shown). The LF filter, HF filter and the decimators together make up a quadrature mirror filter (QMF). Each of the filters can be a finite impulse response filter of conventional form. In use, a line of a field of the input digital video signal is applied, sample-by-sample, to the first stage, to be low pass filtered and high pass filtered by the LF and HF, respectively. Thus, the LF and HF produce outputs comprising low pass filtered and high pass filtered versions of the input line, respectively, the outputs representing the spatial frequency content of the line in the lower and upper halves of the horizontal spatial frequency range. That is, the first stage divides the input line into two sub-bands in the horizontal direction. The decimators decimate (sub-sample) the respective outputs by a factor of two, whereby the total number of samples output by the decimators (together) is the same as the total number of samples in the line.

The second stage is of similar construction to the first stage, except that there are two QMFs each as in the first stage and the

output from each of the decimators of the first stage is passed as an input to a respective one of the two QMFs. Thus, the second stage produces four outputs representing the spatial frequency content of the line in four equal quarters of the horizontal spatial frequency range. That is, the second stage further divides the two sub-bands, into which the input line was divided in the first stage, into four sub-bands in the horizontal direction. The four decimators of the second stage decimate (sub-sample) the respective outputs by a factor of two, whereby the total number of samples output by the decimators of the second stage (together) is the same as the total number of samples in the line.

The third stage is of similar construction to the first stage, except that there are four QMFs each as in the first stage and the output from each of the four decimators of the second stage is passed as an input to a respective one of the four QMFs. Thus, the third stage produces eight outputs representing the spatial frequency content of the line in eight equal one-eighths of the horizontal spatial frequency range. That is, the third stage divides the four sub-bands into which the input line was previously divided into the required eight sub-bands in the horizontal direction. The eight decimators of the third stage decimate (sub-sample) the respective outputs by a factor of two, whereby the total number of samples output by the decimators of the third stage (together) is the same as the total number of samples in the line.

The eight outputs of the third stage, that is of the horizontal filter arrangement 46, are passed to the intermediate field store 48 and stored at positions corresponding to respective one-eighths of a first line thereof. The above process of horizontal filtering is then repeated for all the other lines of the field of the input digital video signal. This results in the intermediate field store 48 containing a version of the field of the input digital video signal that has been filtered into eight sub-bands in the horizontal direction (only). Each line of the field stored in the intermediate field store 48 is divided into eight portions each containing the horizontal spatial frequency information in a respective one of eight sub-bands of the horizontal spatial frequency range of the image that the original field represented. Thus, the horizontally filtered field stored in the

intermediate field store 48 can be considered to be divided into eight columns.

5 The horizontally filtered field stored in the intermediate field store 48 is then fed (under the control of the transpose sequencer 50) into the vertical filter arrangement 52, in which it is filtered into eight sub-bands in the vertical direction in similar manner to that in which filtering into eight sub-bands in the horizontal direction was achieved in the horizontal filter arrangement 46. The horizontally and vertically filtered field is fed on a line-by-line basis into the
10 output field store 54 to be passed from there to the quantiser 14. The store 54 can be considered to have been partitioned into an array of 64 (8 x 8) storage regions, in each of which a respective one of the 64 sub-bands is stored. Thus, successive fields of the input digital video signal are sub-band filtered and passed, duly filtered, to the
15 quantiser 14 after a delay of two field intervals.

The transpose sequencer 50 produces read addresses for the intermediate field store 48, to control reading of the contents thereof into the vertical filter arrangement 52, as follows. As will be recalled, the signal as stored in the intermediate field store 48
20 comprises the lines of the original field, each divided horizontally into eight sub-bands. That is, the signal as stored in the intermediate field store 48 can, as mentioned above, be considered to comprise eight columns. To enable the signal stored in the intermediate field store 48 to be vertically filtered by hardware of
25 the same construction (the vertical filter arrangement 52) used to horizontally filter it, it must be transposed, that is rotated through 90 degrees, as it is read to the vertical filter arrangement 52, so that it comprises eight rows (as opposed to columns). The transpose sequencer 50 addresses the intermediate field store 48 in such a manner
30 as to accomplish this.

The nature of the filtering produced by the combination of the horizontal filter arrangement 46 and the vertical filter arrangement 52 is such that data stored in the output field store 54 is somewhat scrambled and reordered by the output sequencer 56 before being passed
35 to the rest of the apparatus for processing.

Figure 5 illustrates the various sub-band components produced if the input video signal is decomposed both horizontally and vertically

into eight frequency components (this can be considered to be the data stored in the output field store 54 after it has had the reordering of the output sequencer 56 applied to it). Each of these sub-bands or sub-pictures comprises an array of data elements and is represented by one of the blocks in Figure 5. The upper left hand block represents the dc sub-band. This is the band of lowest horizontal and vertical frequency although in practice it doesn't necessarily represent only the constant portions of the signal with strictly zero frequency. This dc sub-band will contain the majority of the dc information of the original input video signal. The relative importance of the remaining sub-bands to the eventual perception of the picture by a viewer varies. Generally speaking, the higher frequency sub-bands are less important to the eventual perception of a viewer. In relation to Figure 5, the frequency which a particular sub-band component represents increases in moving downward and/or rightward in the array of blocks.

Figure 6 illustrates the human psychovisual response to image components of differing spatial frequency. As can be seen, the level of human perception first rises and then steadily diminishes with increasing spatial frequency. This characteristic can be exploited in a data compression system with the realisation that components of higher spatial frequency can be subjected to higher degrees of quantisation with subsequent loss in information, without significantly degrading the perceived image that can be reconstructed.

Figure 7 shows a quantisation matrix that can be applied to the differing sub bands of Figure 5 by the quantiser 14 of Figure 1. The way in which the quantiser operates is that the value of each element within each sub band block is divided by a corresponding quantisation value (α_{00} to α_{77}) and then the integer value of the result is taken. The quantisation value is the product of the entry in the quantisation matrix and a quantisation scaling factor (Q_s) that will be discussed in more detail below. Thus, a value read from the dc luminance sub-band in the top left hand corner would be divided by the corresponding quantisation value of $Q_s \times \alpha_{00}$ and the integer value taken. Similarly, a value read from another sub-band would be divided by the quantisation value $Q_s \times \alpha_{mn}$ appropriate to that sub-band and the integer value taken.

Generally the lowest quantisation values occur in the bands immediately below and to the right of the dc luminance sub-band. This

is because the human psychovisual system is most responsive to these sub bands. The values for the quantisation matrix can be determined by a process of trial and error with subjective viewing tests to see which values give the best perceived image. Alternatively, values can be derived by extending the curve of Figure 6 into three dimensional form to produce a curved surface. The surface is the locus of the curve of Figure 6 rotated through 90° about the sensitivity axis. The sub band frequency blocks lie in different regions of the spatial frequency plane beneath this surface, and by integrating under the surface to find the volume between each frequency block and the surface then a relative quantisation value can be determined.

Figure 8 is a schematic diagram showing the principles of operation of a data compression apparatus providing separate dynamic control of the luminance and chrominance quantisation. A luminance (Y) component data input 60 is supplied to a decorrelator 12' and from there to a luminance quantiser 64. The output of the luminance quantiser is fed via a data sequencer 18' to an entropy encoder 20' which produces a luminance (Y) compressed data output 68. In a similar manner, a chrominance (U/V) component data input 62 is supplied to a decorrelator 12'' and from there to a chrominance quantiser 66. The output from the chrominance quantiser 66 is fed via a data sequencer 18'' to an entropy encoder 20'' which produces a chrominance (U/V) compressed data output 70. The luminance quantiser 64 and the chrominance quantiser 66 operate in a similar manner to the quantiser 14 described earlier. Compression is thus performed separately on the luminance and chrominance component data inputs. The respective luminance and chrominance compressed data outputs 68, 70 may be combined by a combiner (not shown) into a single data stream for further processing.

A combined bit counter 72 detects the quantity of compressed data at the luminance compressed data output 68 and the chrominance compressed data output 70 and produces a feedback signal 74 indicative of the sum of these two quantities. The feedback signal 74 is passed to a quantisation controller 76 which effects separate changes on the degree of luminance quantisation performed by the luminance quantiser 64 and the degree of chrominance quantisation performed by the chrominance quantiser 66 in order that the combined data rate of the

luminance and chrominance compressed data outputs equals or approximates to a desired value. The quantisation controller 76, and the means by which the quantisation controller effects changes to the degree of quantisation performed by the luminance quantiser 64 and the chrominance quantiser 66, will be described in more detail below.

Figure 9 is a graph illustrating the relative changes made by the quantisation controller 76 to the quantisation scaling factors Q_s used in the luminance quantiser 64 and the chrominance quantiser 66.

As mentioned earlier in relation to Figure 7, the degree of quantisation applied to each sub band varies under control of the quantisation matrix. The overall degree of quantisation of an entire block of compressed data (i.e. a collection of 64 sub bands) is varied in dependence upon the quantisation scale factor (Q_s). Each data value being quantised is divided by both the appropriate value from the quantisation matrix and the value of Q_s . If Q_s is high, then the overall degree of quantisation is increased, the degree of compression is increased and the resulting block of compressed data is smaller. If Q_s is low, then the overall degree of quantisation is decreased, the degree of compression is decreased and the resulting block of compressed data is larger.

In Figure 9 two curves are shown, namely a luminance curve (drawn in solid line) and a chrominance curve (drawn in broken line). The vertical axis represents the value of Q_s and the horizontal axis schematically represents the corrective action taken by the quantisation controller 76 in order to change the overall compression factor. Moving to the right on the horizontal axis will tend to increase the overall compression factor, while moving to the left on the horizontal axis will tend to decrease the overall compression factor. At any time the values of Q_s used by the luminance quantiser 64 and the chrominance quantiser 66 will correspond to the position on the vertical axis of respective points on the luminance and chrominance curves, the two points being at the same position on the horizontal axis.

The common origin 78 of the luminance and chrominance curves corresponds to a value of $Q_s = 1$. Moving to the right from the origin 78 it can be seen that Q_s increases at the same rate for both luminance and chrominance compression. At a position 80 the luminance and

chrominance curves diverge so that Q_s for chrominance compression increases more rapidly than Q_s for luminance compression. This divergence continues until a position 82; from the position 82 to a position 84 the value of Q_s continues to increase for luminance compression but remains constant for chrominance compression. From the position 84 the two curves increase at the same rate once again.

It will be seen from the graph in Figure 9 that the overall compression factor is controlled by varying Q_s for luminance compression (the luminance compression factor), Q_s for chrominance compression (the chrominance compression factor) and the ratio of Q_s for luminance compression to Q_s for chrominance compression. It will be appreciated that once two of these parameters have been fixed, the third is also fixed.

Figure 10 illustrates in more detail a data compression apparatus operating with separate dynamic control of the degree of luminance and chrominance quantisation. The apparatus comprises a luminance compression channel 86 (indicated by a broken line boundary), a corresponding chrominance compression channel (not shown), a compressed bit counter 72 and a combiner and header inserter 88. Luminance and chrominance input data are received as blocks of a known size and are compressed to form output data blocks.

In Figure 10 the luminance compression channel 86 and the chrominance compression channel operate in a similar manner, although the quantisation values in the respective quantisation matrices may differ and, as described with reference to Figure 9 above, the variation of Q_s differs. However, for the purposes of the description of Figure 10 the remainder of the operation of the luminance and chrominance compression channels is identical, so the apparatus in Figure 10 will be described with reference to the operation of the luminance compression channel 86 and items shared between the luminance and chrominance compression channels only.

The luminance compression channel 86 shown in Figure 10 is operable to perform a test compression of luminance input data to determine what degree of compression will be achieved with the current value of Q_s . Having determined this a corrected value of Q_s can be generated and applied to the same luminance input data to produce the compressed luminance data that will enter the output data stream of the

data compression apparatus.

In operation, luminance input component data is supplied to a decorrelator 12'. The output from the decorrelator 12' is then fed to a test quantiser 90 that applies a test degree of quantisation based upon the Q_s value stored in a previous Q_s latch 100 for the previous block of input image data. The test quantiser 90 operates in the same way as the luminance quantiser 64 described in relation to Figure 8. Similarly, the test data sequencer 92 and the test entropy encoder 94 operate in the same way as the data sequencer 18' and the entropy encoder 20' from Figure 8 so as to produce a genuine measurement of the degree of compression that will be achieved.

The output from the test entropy encoder 94 in the luminance compression channel is fed to a compressed bit counter 72. The compressed bit counter 72 also receives the output of the corresponding test entropy encoder in the chrominance compression channel (not shown). The compressed bit counter 72 detects the quantity of data (bit count) output during test compression of a block of input image data by the test entropy encoders in the luminance compression channel and the chrominance compression channel and produces a feedback signal indicative of the sum of these two quantities. At the end of the test compression the output from the compressed bit counter 72 is fed to a programmable read-only memory (PROM) 98 in the luminance compression channel which calculates a replacement value of Q_s according to the luminance curve shown in Figure 9. Similarly a corresponding PROM 98 in the chrominance compression channel receives the same output from the compressed bit counter 72 and calculates a replacement value of Q_s for use in chrominance compression. The manner in which the PROMs 98 perform these calculations will be described below with reference to Figure 11.

The PROM 98 in the luminance compression channel 86 outputs a corrected value of Q_s to the previous Q_s latch 100 for use in the next test compression and to a Q_s latch 102 for use in compression of the current luminance input data. The output from the Q_s latch 102 is fed to the luminance quantiser 64 where it controls the overall degree of quantisation. Whilst the data from the decorrelator 12' was undergoing test compression it was also being processed by the delay circuit 104 which has the effect of delaying the arrival of that block of data at

the luminance quantiser 64 until the corrected value of Q_s has been stored within the Q_s latch 102. In this way, a corrected value of Q_s can be applied to the same block of data upon which its correction was based.

5 The output from the luminance quantiser 64 is entropy encoded by an entropy encoder 20' coupled to a field store 60 which stores the compressed data prior to its final formatting by the entropy encoder 20' and output from the system.

10 It will be appreciated that when decompressing a block of compressed data that has had a variable degree of quantisation applied to it, the decompression apparatus should be responsive to the degree of quantisation that was applied to that block of data. To this end, the value stored within the Q_s latch 102 is supplied as header data for the block of compressed data stored within the field store 60. The
15 header inserter and combiner unit 88 prefixes this header data to the compressed block of data. This enables a decompression apparatus to read from the header data what quantisation scaling factor was applied to each block of compressed data. The combination of the respective PROMs 98 and Q_s latches 102 in the luminance and chrominance channels
20 can be thought of as a quantisation controller 76.

 The output from the entropy encoder 20' is also fed to a second compressed bit counter 108, the output of which is continuously fed to an overflow detector 106. The overflow detector 106 compares the instantaneous compressed bit count in each compression channel with the
25 maximum predetermined size of the block of compressed data being compiled and outputs an overflow indicator when the predetermined maximum size has been exceeded. The overflow indicator triggers the entropy encoder 20' to generate a code not normally produced for compressed data, so that the overflow indicator will not be confused
30 with genuine compressed data. An appropriate code can be chosen by reference to the entropy encoder code table.

 The compressed data may then, for example, be recorded as individual tracks on a magnetic tape for subsequent recovery and decompression. It will be understood that the blocks of compressed
35 data could be handled in many other different manners before decompression.

 Figure 11 illustrates one technique whereby the PROMs 98 in the

luminance compression channel and the chrominance compression channel determines what change in Q_s to make between different blocks of input data. Each PROM 98 reads the output from the compressed bit counter 72 and subtracts from this the known predetermined maximum size of the block of compressed data. This yields a difference value D. The PROM 98 then refers to a look-up table in which data representing the graph of Figure 11 is stored. A change in quantisation appropriate to the measured difference value D and the present value of Q_s is read from this look-up table. The change is in accordance with the appropriate curve in Figure 9. If the difference value is positive, then the block of compressed data is too large and a positive value is added the current value Q_s to produce a new value of Q_s for use on the next block of data. A larger value of Q_s gives more compression. If the difference value is negative, then the block of compressed data is too small and a negative value is added to the current value of Q_s . The new value of Q_s is then fed to the Q_s latch 66.

This technique copes with both overruns and underruns in compressed block size. If the quantisation applied has been too severe, then the compressed block will be too small and the PROM 98 will cause a decrease in the value of Q_s in accordance with the appropriate curve shown in Figure 9. If the quantisation has not been severe enough, then the compressed block will be too large and the PROM 98 will cause an increase in the value of Q_s . In this way, the size of the compressed blocks is held to be substantially equal to the predetermined maximum size. This helps to ensure that overruns do not occur, whilst not applying too severe a degree of quantisation.

Figure 12 illustrates a data decompression apparatus complementary to the data compression apparatus shown in Figure 10. The decompression apparatus comprises a header reader and data splitter 110, a luminance decompression channel 112 and a chrominance decompression channel 114. The luminance and chrominance decompression channels are similar in structure and operation, so only the luminance decompression channel will be described in detail. The luminance decompression channel comprises an entropy decoder 116 and field store 118, a data resequencer 120, a luminance dequantiser 122 and an interpolator 124.

The compressed data from the data compression apparatus is

transferred to the data decompression apparatus via one of the mechanisms previously described (such as by magnetic recording and subsequent reproduction) and is then passed through the header reader and data splitter 110. The luminance compressed data then undergoes reformatting by the entropy decoder 116 and storage in the field store 118. The header reader and data splitter 110 reads the value of Q_s used for luminance compression and put into the header data of the compressed block, and stores this within Q_s latch 84. The compressed luminance data itself then streams through the header reader 80 through the entropy decoder 116 into the field store 118.

If the entropy decoder 116 detects the occurrence of an overflow indicator code within the data stream it recognises that an overflow occurred when that data was compressed and sets all the output data values not determinable from the compressed data to a predetermined value within the field store 118, e.g. sets all the undeterminable values to zero.

The Q_s latch 84 supplies the value of Q_s to the luminance dequantiser 122 during dequantisation of compressed luminance data from each block. Accordingly, the compressed luminance data has a degree of dequantisation applied to it that matches the degree of quantisation applied during compression.

The operation of the entropy decoder 116, the data resequencer 120 and the interpolator 124 are complementary to the operation of the entropy encoder 20, the data sequencer 18 and the decorrelator 12. The luminance dequantiser 122 multiplies each compressed data value by the appropriate value from the quantisation matrix and the value of Q_s from the Q_s latch 126, to effect the complementary action to that carried out by the luminance quantiser 64.

The chrominance decompression channel 114 operates in a similar manner. The header reader and data splitter 110 supplies the quantisation scaling factor Q_s used for chrominance compression and stored in the header data to a chrominance (U/V) Q_s latch 128 and supplies compressed chrominance data to the entropy decoder and overflow detector 116 in the chrominance decompression channel 114. Dequantisation is performed by a chrominance dequantiser 130.

CLAIMS

1. Data compression apparatus for compressing luminance and chrominance portions of input image data to form compressed image data having an overall compression factor, the data compression apparatus comprising:

luminance compression means for compressing the luminance portion of the input image data by a luminance compression factor;

chrominance compression means for compressing the chrominance portion of the input image data by a chrominance compression factor;

feedback means for supplying a feedback signal dependent on the overall compression factor; and

control means responsive to the feedback signal and coupled to the luminance compression means and the chrominance compression means for controlling the overall compression factor by varying the luminance compression factor, the chrominance compression factor and the ratio of the luminance compression factor to the chrominance compression factor.

2. Apparatus according to claim 1, in which the control means varies the luminance compression factor to be less than or equal to the chrominance compression factor.

3. Apparatus according to claim 2, in which the control means varies the luminance compression factor to be equal to the chrominance compression factor for one or more predetermined ranges of values of the overall compression factor and to be less than the chrominance compression factor for values of the overall compression factor which are not in the one or more predetermined ranges of values.

4. Apparatus according to any one of the preceding claims, in which the luminance compression means and the chrominance compression means each comprise:

a decorrelator for transforming the image data into transformed data representing a plurality of frequency components in the spatial frequency domain;

a quantiser for applying a variable degree of quantisation to the transformed data; and

an entropy encoder for encoding the quantised transformed data.

5 5. Apparatus as claimed in any preceding claim, further comprising means for associating a data header with the compressed image data indicative of the luminance compression factor and the chrominance compression factor corresponding to that compressed image data.

10 6. Apparatus according to any one of the preceding claims, in which blocks of the input image data are compressed to form blocks of compressed data, and in which the feedback means comprises a compressed bit counter for storing a count of bits in blocks of compressed data.

15 7. Data decompression apparatus for decompressing compressed image data having luminance and chrominance portions, the data decompression apparatus comprising:

20 header reading means for reading a data header associated with the compressed image data and providing separate indications of the degree of luminance compression and the degree of chrominance compression applied to the luminance and the chrominance portions respectively of the compressed image data;

 luminance decompression means, responsive to the header reading means, for decompressing the luminance portion of the compressed image data by a degree dependent on the degree of luminance compression; and

25 chrominance decompression means, responsive to the header reading means, for decompressing the chrominance portion of the compressed image data by a degree dependent on the degree of chrominance compression.

30 8. A method of compressing luminance and chrominance portions of input image data to form compressed image data having an overall compression factor, the method comprising the steps of:

 compressing the luminance portion of the input image data by a luminance compression factor;

35 compressing the chrominance portion of the input image data by a chrominance compression factor;

 varying the luminance compression factor, the chrominance compression factor and the ratio of the luminance compression factor to

the chrominance compression factor in order to control the overall compression factor.

- 5 9. A method of decompressing compressed image data having luminance and chrominance portions, the method comprising the steps of:
- reading a data header associated with the compressed image data and providing separate indications of the degree of luminance compression and the degree of chrominance compression applied to the luminance and the chrominance portions respectively of the compressed
- 10 image data;
- decompressing the luminance portion of the compressed image data by a degree dependent on the degree of luminance compression; and
- decompressing the chrominance portion of the compressed image data by a degree dependent on the degree of chrominance compression.
- 15
10. Image data storage apparatus comprising image data compression apparatus as claimed in any one of claims 1 to 6.
- 20 11. Image data storage apparatus comprising image data decompression apparatus as claimed in claim 7.
12. Data compression apparatus substantially as hereinbefore described with reference to the accompanying drawings.
- 25 13. Data decompression apparatus substantially as hereinbefore described with reference to the accompanying drawings.
14. A method of data compression substantially as hereinbefore described with reference to the accompanying drawings.
- 30 15. A method of data decompression substantially as hereinbefore described with reference to the accompanying drawings.

Patents Act 1977

Examiner's report to the Comptroller under
Section 17 (The Search Report)

Application number

9204329.8

Relevant Technical fields

(i) UK Cl (Edition K) H4F [FRT, FRC, FRP, FRG, FRR, FRD,
FRM, FRS, FRX]

Search Examiner

R KING

(ii) Int CL (Edition 5) H04N

Databases (see over)

(i) UK Patent Office

Date of Search

30 JUNE 1992

(ii) ONLINE DATABASE: WPI

Documents considered relevant following a search in respect of claims

1-6

Category (see over)	Identity of document and relevant passages	Relevant to claim(s)
	NONE	1

Category	Identity of document and relevant passages	Relevant to claim(s)

Categories of documents

X: Document indicating lack of novelty or of inventive step.

Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.

A: Document indicating technological background and/or state of the art.

P: Document published on or after the declared priority date but before the filing date of the present application.

E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.

&: Member of the same patent family, corresponding document.

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).